arXiv:1603.06850v3 [cs.FL] 12 May 2016

Array Folds Logic*

Przemysław Daca, Thomas A. Henzinger, and Andrey Kupriyanov

IST Austria, Austria

Abstract. We present an extension to the quantifier-free theory of integer arrays which allows us to express counting. The properties expressible in Array Folds Logic (AFL) include statements such as "the first array cell contains the array length," and "the array contains equally many minimal and maximal elements." These properties cannot be expressed in quantified fragments of the theory of arrays, nor in the theory of concatenation. Using reduction to counter machines, we show that the satisfiability problem of AFL is PSPACE-complete, and with a natural restriction the complexity decreases to NP. We also show that adding either universal quantifiers or concatenation leads to undecidability.
AFL contains terms that fold a function over an array. We demonstrate that folding, a well-known concept from functional languages, allows us to concisely summarize loops that count over arrays, which occurs frequently in real-life programs. We provide a tool that can discharge proof obligations in AFL, and we demonstrate on practical examples that our decision procedure can solve a broad range of problems in symbolic testing and program verification.

1 Introduction

Arrays and lists (or, more generally, sequences) are fundamental data structures both for imperative and functional programs: hardly any real-life program can work without processing sequentially-ordered data. Testing and verification of array- and list-manipulating programs is thus a task of crucial importance. Almost any non-trivial property about these data structures requires some sort of universal quantification; unfortunately, the full first-order theories of arrays and lists are undecidable. This has motivated researches to investigate fragments with restricted quantifier prefixes, and has given rise to numerous logics that can describe interesting properties of sequences, such as partitioning or sortedness. These logics have efficient decision procedures and have been successfully applied to verify some important aspects of programs working with arrays and lists: for example, the correctness of sorting algorithms.

However, an important class of properties, namely, *counting* over arrays, has eluded researchers' attention so far. In addition to the examples from the abstract, this includes statements such as "the histogram of the input data satisfies

* This research was supported in part by the European Research Council (ERC) under grant 267989 (QUAREM) and by the Austrian Science Fund (FWF) under grants S11402-N23 (RiSE and SHiNE) and Z211-N23 (Wittgenstein Award).

```
min = max = a[0];
j = k = 0;
for(i=0;i<size(a);i++) {
  if(a[i]<min) { min=a[i]; j=1; }
  if(a[i]==min) j++;
}
for(i=0;i<size(a);i++) {
  if(a[i]>max) { max=a[i]; k=1; }
  if(a[i]==max) k++;
}
assert(j==k);
```

(a) `C` language

$$
\begin{aligned}
&\exists\, min, max, i_1, i_2, j, k\ .\\
&0 \le i_1 < |a| \ \wedge\ 0 \le i_2 < |a| \ \wedge\\
&a[i_1] = min \ \wedge\ a[i_2] = max \ \wedge\\
&\forall i.(a[i] \ge min) \ \wedge\\
&\forall i.(a[i] \le max) \ \wedge\\
&j = \left|\{i \mid a[i] = min\}\right| \ \wedge\\
&k = \left|\{i \mid a[i] = max\}\right| \ \wedge\\
&j = k
\end{aligned}
$$

(b) Quantified arrays + cardinality

$$
\begin{aligned}
&0 \le i_1 < |a| \ \wedge\ 0 \le i_2 < |a| \ \wedge\ a[i_1] = min \ \wedge\ a[i_2] = max \ \wedge\\
&fold_a \tbinom{0}{0} \tbinom{\mathbf{e}=min \Rightarrow \mathbf{c_1{+}{+}}}{\mathbf{e}>min \Rightarrow skip} = \tbinom{|a|}{j} \ \wedge\ fold_a \tbinom{0}{0} \tbinom{\mathbf{e}=max \Rightarrow \mathbf{c_1{+}{+}}}{\mathbf{e}<max \Rightarrow skip} = \tbinom{|a|}{k} \ \wedge\ j = k
\end{aligned}
$$

(a) Array Folds Logic

Fig. 3: A toy array problem

the given distribution," or "the packet adheres to the requirements of the given type-length-value (TLV) encoding (e.g., of the IPv6 options)." Such properties, though crucial for many applications, cannot be expressed in decidable fragments of the first-order theory of arrays, nor in the decidable extensions of the theory of concatenation.

In this paper we present *Array Folds Logic* (*AFL*), which is an extension of the quantifier-free theory of integer arrays. But instead of introducing quantifiers, we introduce counting in the form of *fold* terms. Folding is a well-known concept in functional languages: as the name suggests, it folds some function over an array, i.e., applies it to every element of the array in sequence, while preserving the intermediate result.

To illustrate the kind of problems we are dealing with, consider the following toy example: given an array, accept it if the number of minimum elements in the array is the same as the number of maximum elements in the array. E.g., the array $[1, 2, 7, 4, 1, 3, 7, 5]$ is accepted (because there are two 1's and two 7's), while the array $[1, 2, 7, 4, 1, 3, 6, 5]$ is rejected (because there is only one 7).

Written in a programming language like `C`, the problem can be solved by the piece of code shown in Figure 1a, but such *explicit* solution cannot express verification conditions for *symbolic* verification and testing. We can use the quantified theory of arrays mixed with assertions about cardinality of sets, as in Figure 1b. Unfortunately, such a combination is undecidable (by a reduction from Hilbert's Tenth Problem: replace *fold*s with cardinalities in the proof of Theorem 2).

The solution we propose is shown in Figure 2a: in the example formula, the first *fold* applies a function to array a. The vector in the first parentheses gives initial values for the array index and counter $\mathbf{c}_1$; the function is folded over the array starting from the initial index. Index variable $\mathbf{i}$ is implicit, and it is incremented at each iteration. The function itself is given in the second parentheses, and has two branches. The first branch *counts* the number of positions with elements equal to *min* in counter $\mathbf{c}_1$. The second branch *skip*s when the current array element $\mathbf{e}$ is greater than the (guessed, existentially quantified) variable *min*. When $\mathbf{e} < min$, the implicit *break* statement is executed, and the *fold* terminates prematurely. The result of the *fold* is compared to the vector which asserts that the final value of the array index equals to the array size $|a|$ (which means no *break* was executed), and the final value of $\mathbf{c}_1$ equals to j. The positions where elements are equal to *max*, are counted in the second *fold*, and the equality between these two counts is asserted. The ability to count over arrays with unbounded elements is a unique feature of Array Folds Logic.

This paper makes the following contributions:

1. We define a new logic, called AFL, that can express interesting and non-trivial properties of counting over arrays, which are orthogonal to the properties expressible by other logics. Additionally, AFL can concisely summarize loops with internal branching that traverse arrays and perform counting, enabling verification and symbolic testing of programs with such loops.

2. We show that the satisfiability problem for AFL is **PSPACE**-complete, and with a natural restriction the complexity decreases to **NP**. We provide a decision procedure for AFL, which works by a reduction to the emptiness of (symbolic) reversal-bounded counter machines, which in turn reduces to the satisfiability of existential Presburger formulas. We show that adding either universal quantifiers or concatenation leads to undecidability.

3. We implemented tool AFolder [13] that can discharge proof obligations in AFL, and we demonstrate on real-life examples that our decision procedure can solve a broad range of problems in symbolic testing and program verification.

Related work. Our logic is related to the quantified fragments of the theory of arrays such as [10][20][21][4]. These logics allow restricted quantifier prefixes, and their decision procedures work by rewriting to the (parametric) theories of array indices and elements (Presburger arithmetic being the most common case) [10][4], or by reduction to flat counter automata with difference bound constraints [20][21]. An interesting alternative is provided in [36], where the quantification is arbitrary, but array elements must be bounded by a constant given a priori; the decision procedure works by a reduction to WS1S. A separate line of work is presented by the theory combination frameworks of [17][15], where the quantifier-free theory of arrays is extended by *injective* predicate and *domain* function [17], or with *map* and *constant-value* combinators [15]. The theory of concatenation and its extensions [11][16][28] are also related; their decision procedures work by reduction to Makanin's algorithm for solving word equations [29]. AFL can express some properties that are also expressible in these logics, such as *boundedness*, *partitioning*, or *periodicity*; other properties,

such as *sortedness*, are not expressible in AFL. The counting properties that constitute the core of AFL are not expressible in any of the above logics. We compare the expressive power of AFL and other logics in Section 2.3.

There are numerous works on loop acceleration and summarization [12][8][26], also in the context of verification and symbolic testing [33][9][18][23] and array-manipulating programs [7][3][5]. Our logic allows one to summarize loops with internal branching and counting, which are outside of the scope of these works.

The decision procedure for AFL is based on decidability results for emptiness of reversal-bounded counter machines [24][19][25], on the encoding of this problem into Presburger arithmetic [22], and on the computation of Parikh images for NFAs [34]. In Section 5 we extend the encoding procedure to symbolic counter machines, and present some substantial improvements that make it efficient for solving practical AFL problems.

2 Array Folds Logic

We assume familiarity with the standard syntax and terminology of many-sorted first-order logics. We use vector notation: $\boldsymbol{v} = (v_1, \ldots, v_n)$ denotes an ordered sequence of terms. For two vectors $\boldsymbol{u}$ and $\boldsymbol{v}$, we write their concatenation as $\boldsymbol{uv}$.

Within this paper we consider the domains of arrays, array indices, and array elements to be $\mathbb{A} = \mathbb{Z}^*$, $\mathbb{N} = \{ 0, 1, \ldots \}$, and $\mathbb{Z} = \{ \ldots, -1, 0, 1, \ldots \}$ respectively.

Presburger arithmetic has the signature $\Sigma_{\mathbb{Z}} = \{ 0, 1, +, < \}$; we use it for array indices and elements, as well as other arithmetic assertions, possibly with embedded array terms. We write *true* and *false* to denote a valid and an unsatisfiable Presburger formula, respectively.

The theory of integer-indexed arrays extends Presburger arithmetic with functions *read*, and *write*, and has the signature $\Sigma_A = \Sigma_{\mathbb{Z}} \cup \{ \cdot[\cdot], \cdot\{\cdot \leftarrow \cdot\} \}$. The *read* function $a[i]$ returns the i-th element of array a, and the *write* function $a\{i \leftarrow x\}$ returns array a where the i-th element is replaced by x. These functions should satisfy the *read*-over-*write* axioms as described by McCarthy [31]; the decision procedure for the quantifier-free array theory is presented in [35].

2.1 Syntax

Array Folds Logic (AFL) extends the quantifier-free theory of integer arrays with the ability to perform counting. The extension works by incorporating *fold* terms into arithmetic expressions; such a term folds some function over the array by applying it to each array element consecutively.

AFL contains the following sorts: array sort ASort, integer sort ISort, Boolean sort BSort, and two enumerable sets of sorts for integer vectors VSort^m and functional constants $\mathsf{FSort}^m = \mathsf{VSort}^m \times \mathsf{ISort} \to \mathsf{VSort}^m$, for each $m \in \mathbb{N}$, $m > 0$. The syntax of the AFL terms is shown in Table 1; a and b denote array variables, x denotes an integer variable, n and m denote integer constants.

Array terms A of sort ASort are represented either by an array variable a, or by the *write* term $a\{T \leftarrow T\}$.

Integer terms T of sort ISort can be integer constants $n \in \mathbb{Z}$, integer variables x, integer addition, *read* term $a[T]$ for the index represented as an integer term, or the term $|a|$, which represents the length of array a.

Boolean terms B of sort BSort are formed by array equality, usual Presburger and Boolean operators, and equality between vectors of sort VSort^m.

Vector terms V^m of sort VSort^m are either a list of m integer terms, or a *fold* term. The former is written as a vertical list in parentheses; they can be omitted when $m = 1$. The latter, written as $\mathit{fold}_a\ v\ f$, represents the result of the transformation of an input vector v of sort VSort^m by folding a functional constant f of sort FSort^m over an array a. The first element of v specifies an initial value of the array index; the remaining elements give initial values for the counters that can be used inside f. The resulting vector after the transformation gives the final values for the array index and the counters.

Functional constants (when no confusion can arise, we call them *functions*) F^m of sort FSort^m can only be a parenthesized list of branches (guarded commands); the length of the list is unrelated to m. A function f of sort FSort^m can refer to the following implicitly declared variables: $\mathbf{e}$ for the currently inspected array element; $\mathbf{i}$ for the current array index; $\mathbf{c}_1, \ldots, \mathbf{c}_{m-1}$ for the counters; $\mathbf{s}$ for the state (control flow) variable. All other variables that occur inside f are considered as free variables of sort ISort.

Guards are conjunctions of *atomic guards*, which can compare array elements, indices, and counters to integer terms; the state variable can only be compared to integer constants. *Updates* are lists of *atomic updates*; they can increment or decrease counters by a constant, assign a constant to the state variable, *skip*, i.e. perform no updates, or execute a *break* statement, which terminates the *fold* at the current position. Counter or state updates define a function $\mathbb{Z} \to \mathbb{Z}$. Guards and updates translate into logical formulas that either constraint the current variable values, or relate the current and the next-state (primed) variable values in the obvious way; we denote this translation by Φ. E.g., the update $\mathit{upd} \equiv (\mathbf{c}_1\ \texttt{+=}\ n)$ defines the formula $\Phi(\mathit{upd}) \equiv (\mathbf{c}_1' = \mathbf{c}_1 + n)$.

We require that guards of all branches are mutually exclusive. There is an implicit "catch-all" branch with the *break* statement, whose guard evaluates to *true* exactly when guards of all other branches evaluate to *false*. We also require that each branch contains at most one update for each implicit variable.

We restrict the control flow in functions, which is defined by state variable $\mathbf{s}$. Notice that $\mathbf{s}$ is syntactically finite state. Thus, given a set of function branches Br, we define an edge-labeled control flow graph $G = \langle S, E, \gamma \rangle$, where:

- states $S = \{0\} \cup \{\ n \mid \mathbf{s} \leftarrow n \in Br\ \}$;
- edges $E = \bigcup_{\mathit{grd} \Rightarrow \mathit{upd}\ \in\ Br} \{\ (s_1, s_2) \mid s_1 \models \mathit{grd} \wedge s_2 = \mathit{ite}(\mathbf{s} \leftarrow n \in \mathit{upd}, n, s_1)\ \}$;
- γ is the labeling of edges with the set of formulas $\Phi(\mathit{grd})$ and $\Phi(\mathit{upd})$ for each guard or update which occurs in the same branch.

We require that edges in the strongly-connected components of G are labeled with counter updates that are, for each counter, all non-decreasing, or all non-

$$
\begin{array}{lcl}
A & ::= & a \mid a\{T \leftarrow T\} \\
T & ::= & n \mid x \mid T + T \mid a[T] \mid |a| \\
B & ::= & a = b \mid T = T \mid T < T \mid \neg B \mid B \wedge B \mid V^m = V^m \\
V^m & ::= & \binom{T}{\cdots}{}_{T} \mid \mathit{fold}_a\, V^m\, F^m \\
F^m & ::= & \binom{\mathit{grd} \Rightarrow \mathit{upd}}{\cdots}{}_{\mathit{grd} \Rightarrow \mathit{upd}} \\
\mathit{grd} & ::= & \mathbf{e} \approx T \mid \mathbf{i} \approx T \mid \mathbf{c}_m \approx T \mid \mathbf{s} \approx n \mid \mathit{grd} \wedge \mathit{grd} \qquad (\approx \in \{>, <, =, \neq\}) \\
\mathit{upd} & ::= & \mathbf{c}_m \texttt{+=} n \mid \mathbf{s} \leftarrow n \mid \mathit{skip} \mid \mathit{break} \mid \mathit{upd}\,;\,\mathit{upd}
\end{array}
$$

Table 1: Syntax of AFL.

increasing. Thus, G is a DAG of SCCs, where counters within each SCC behave in a monotonic way. We use this restriction to derive from f a reversal-bounded counter machine (see Definition 2).

The presented syntax is minimal and can be extended with convenience functions and predicates such as $\{-, n\cdot, , \leq, \geq, \vee, \texttt{++}, \texttt{--}, \texttt{-=}n\}$ in the usual way. We allow to use $*$ to denote the absence of constraints: this is useful for vector notation. We replace each $*$ in the formula with a unique unconstrained variable.

2.2 Semantics

For a given AFL formula ϕ, we denote the sets of free variables of ϕ of sort ASort and ISort by Var_A and Var_I, respectively. All free variables are implicitly existentially quantified. For functions of sort FSort^m, we denote by FV^m the set of their implicit variables $\{\mathbf{i}, \mathbf{c}_1, \ldots, \mathbf{c}_{m-1}, \mathbf{s}\}$.

The treatment of array writes and reads in the quantifier-free array theory is standard [35], and we do not elaborate on it here. Array equalities partition the set of array variables into equivalence classes; all other constraints are then translated into constraints over a representative of the corresponding equivalence class.

An *interpretation* for AFL is a tuple $\sigma = \langle \lambda, \mu \rangle$, where $\lambda : \mathit{Var}_I \to \mathbb{Z}$ assigns each integer variable an integer, and $\mu : \mathit{Var}_A \to \mathbb{Z}^*$ assigns each array variable a finite sequence of integers.

The semantics of an AFL term t under the given interpretation σ is defined by the *evaluation* $[t]^\sigma$. Terms that constitute functions are evaluated in the additional *context* κ. For a function f of sort FSort^m, $\kappa : FV^m \to \mathbb{Z}^{m+1}$ maps internal variables of f to integers. The evaluation of Presburger, Boolean, and array terms is standard; the remaining ones are shown in Table 2. We give some explanations here (the remaining semantic rules are self-explanatory):

1. Vector equality resolves to a conjunction of equalities between components.

1.	$\left[\binom{t^1_1}{\cdots}{=}\binom{t^2_1}{\cdots}{t^1_m}{t^2_m}\right]^\sigma$	$\equiv$	$\left(\left[t^1_1\right]^\sigma = \left[t^2_1\right]^\sigma\right) \wedge \ldots \wedge \left(\left[t^1_m\right]^\sigma = \left[t^2_m\right]^\sigma\right)$
2.	$[\mathit{fold}_a\ v\ f]^\sigma$	$\equiv$	$[\mathit{fold}_a\ v\ f]^{\sigma,\kappa}$, where $\kappa(FV^m) = \binom{v}{0}$
3.	$[\mathit{fold}_a\ v\ f]^{\sigma,\kappa}$	$\equiv$	if $([\mathbf{i}]^\kappa < 0)$ or $([\mathbf{i}]^\kappa \geq \lvert a\rvert)$ or $(\mathit{false} \in [f]^{\sigma,\kappa})$ then v else $\left[\mathit{fold}_a\ v'\ f\right]^{\sigma,\kappa'}$, where $\kappa'(FV^m) = \binom{v'}{\mathbf{s}'} = [f]^{\sigma,\kappa}\left(\kappa(FV^m)\right)$
4.	$[f]^{\sigma,\kappa}\binom{v_1}{\cdots}{v_m}$	$\equiv$	$\binom{v'_1}{\cdots}{v'_m}$, where $v'_j \equiv$ if $\mathit{upd}(v_j) \in [f]^{\sigma,\kappa}$ then $\mathit{upd}(v_j)$, else v_j
5.	$\left[\binom{grd_1 \Rightarrow upd_1}{\cdots}{grd_m \Rightarrow upd_m}\right]^{\sigma,\kappa}$	$\equiv$	$\{\mathbf{i}' = \mathbf{i}+1\} \cup [grd_1 \Rightarrow upd_1]^{\sigma,\kappa} \cup \ldots \cup [grd_m \Rightarrow upd_m]^{\sigma,\kappa}$
6.	$[grd \Rightarrow upd]^{\sigma,\kappa}$	$\equiv$	if $[grd]^{\sigma,\kappa} = \mathit{true}$ then $[upd]^{\sigma,\kappa}$ else $\emptyset$
7.	$[\mathbf{e} \approx t]^{\sigma,\kappa}$	$\equiv$	$[\mathbf{e}]^\kappa \approx [t]^\sigma$ (similarly for $\mathbf{i} \approx T$, $\mathbf{c}_m \approx T$, $\mathbf{s} \approx n$)
8.	$[grd_1 \wedge grd_2]^{\sigma,\kappa}$	$\equiv$	$[grd_1]^{\sigma,\kappa} \wedge [grd_2]^{\sigma,\kappa}$
9.	$[upd_1; upd_2]^{\sigma,\kappa}$	$\equiv$	$[upd_1]^{\sigma,\kappa} \cup [upd_2]^{\sigma,\kappa}$
10.	$[\mathbf{c}_m \texttt{ += } n]^{\sigma,\kappa}$	$\equiv$	$\{\mathbf{c}'_m = \mathbf{c}_m + n\}$
11.	$[\mathbf{s} \leftarrow n]^{\sigma,\kappa}$	$\equiv$	$\{\mathbf{s}' = n\}$
12.	$[\mathit{skip}]^{\sigma,\kappa}$	$\equiv$	$\emptyset$
13.	$[\mathit{break}]^{\sigma,\kappa}$	$\equiv$	$\{\mathit{false}\}$

Table 2: Semantics of AFL

2. A *fold* term evaluates in the initial context that is defined by the given initial vector of counters v, and assigns 0 to state variable $\mathbf{s}$.
3. A contextual *fold* term checks whether the array index is out of bounds, or a *break* statement is executed in the current context (this is the only way for $[f]^{\sigma,\kappa}$ to contain *false*). If yes, *fold* terminates, and returns the current vector v. Otherwise *fold* continues with the updated vector and context.
4. If an update $\mathit{upd}(v_j)$ for some variable v_j is present in the function evaluation, then it is applied. Otherwise, the old variable value is preserved.
5. An evaluation of a function, represented by a list of branches, is a union of updates from its branch evaluations. Index $\mathbf{i}$ is always incremented by 1.
6. A guarded command evaluates to its update if its guard evaluates to *true*.
7. A comparison over an internal variable evaluates it in the context κ, and the comparison term is evaluated in the interpretation σ.

2.3 Expressive Power

Here we give some example properties that are expressible in AFL, and compare its expressive power to other decidable array logics.

1. **Boundedness**. All elements of array a belong to the interval $[l, u]$.

$$\mathit{fold}_a\big(0\big)\big(l \leq \mathbf{e} \leq u \Rightarrow \mathit{skip}\big) = |a|$$

2. **Partitioning**. Array a is partitioned if there is a position p such that all elements before p are smaller or equal than all elements at or after p.

$$\mathit{fold}_a\binom{0}{}\binom{\mathbf{i}{<}p \wedge \mathbf{e}{\leq}a[p] \Rightarrow \mathit{skip}}{\mathbf{i}{\geq}p \wedge \mathbf{e}{\geq}a[p] \Rightarrow \mathit{skip}} = |a|$$

3. **Periodicity**. Array a is of the form $(01)^*$:

$$\mathit{fold}_a\binom{0}{}\binom{\mathbf{s}{=}0 \wedge \mathbf{e}{=}0 \Rightarrow \mathbf{s}{\leftarrow}1}{\mathbf{s}{=}1 \wedge \mathbf{e}{=}1 \Rightarrow \mathbf{s}{\leftarrow}0} = |a|$$

4. **Pumping**. Array a is of the form 0^n1^n (a canonical non-regular language; $0^n1^n2^n$, a non-context-free language, is equally expressible):

$$\mathit{fold}_a\binom{0}{0}\binom{\mathbf{s}{=}0 \wedge \mathbf{e}{=}0 \Rightarrow \mathbf{c}_1\mathbf{{+}{+}}}{\mathbf{s}{=}0 \wedge \mathbf{e}{=}1 \Rightarrow \mathbf{c}_2\mathbf{{+}{+}} \wedge \mathbf{s}{\leftarrow}1}{\mathbf{s}{=}1 \wedge \mathbf{e}{=}1 \Rightarrow \mathbf{c}_2\mathbf{{+}{+}}} = \binom{|a|}{n}{n}$$

5. **Equal Count**. Arrays a and b have equal number of elements greater than l:

$$\binom{|a|}{n} = \mathit{fold}_a\binom{0}{0}\binom{\mathbf{e}{>}l \Rightarrow \mathbf{c}_1\mathbf{{+}{+}}}{\mathbf{e}{\leq}l \Rightarrow \mathit{skip}} \wedge \binom{|b|}{n} = \mathit{fold}_b\binom{0}{0}\binom{\mathbf{e}{>}l \Rightarrow \mathbf{c}_1\mathbf{{+}{+}}}{\mathbf{e}{\leq}l \Rightarrow \mathit{skip}}$$

6. **Histogram**. The histogram of the input data in array a satisfies the distribution $H\big(\{i \mid a[i] < 10\}\big) \geq 2H\big(\{i \mid a[i] \geq 10\}\big)$:

$$\mathit{fold}_a\binom{0}{0}\binom{\mathbf{e}{<}10 \Rightarrow \mathbf{c}_1\mathbf{{+}{+}}}{\mathbf{e}{\geq}10 \Rightarrow \mathit{skip}} = \binom{|a|}{h_1} \wedge \mathit{fold}_a\binom{0}{0}\binom{\mathbf{e}{\geq}10 \Rightarrow \mathbf{c}_1\mathbf{{+}{+}}}{\mathbf{e}{<}10 \Rightarrow \mathit{skip}} = \binom{|a|}{h_2} \wedge h_1 \geq 2h_2$$

7. **Length of Format Fields**. The array contains two variable-length fields. The first two elements of the array define the length of each field; they are followed by the fields themselves, separated by 0:

$$\mathit{len}_1 = a[0] \wedge \mathit{len}_2 = a[1] \wedge \mathit{fold}_a\binom{2}{0}{0}\binom{\mathbf{s}{=}0 \wedge \mathbf{e}{\neq}0 \Rightarrow \mathbf{c}_1\mathbf{{+}{+}}}{\mathbf{s}{=}0 \wedge \mathbf{e}{=}0 \Rightarrow \mathbf{s}{\leftarrow}1}{\mathbf{s}{=}1 \wedge \mathbf{e}{\neq}0 \Rightarrow \mathbf{c}_2\mathbf{{+}{+}}} = \binom{|a|}{\mathit{len}_1}{\mathit{len}_2}$$

Comparison with other logics. Most decidable array logics can specify universal properties over a *single* index variable like (1) above; AFL uses *fold*s to express such universal quantification. Properties that require universal quantification over *several* index variables, like sortedness, are inexpressible in AFL (it can simulate some of such properties, like partitioning (2), using a combination of *fold*s with existential guessing). Periodic facts like (3) are inexpressible in [10], but AFL as well as [20][16] can express it. Counting properties such as (4)–(7), which constitute the core of AFL, are not expressible in other decidable logics over arrays and sequences.

3 Motivating Example

As a motivating example to illustrate applications of our logic, we consider a parser for the Markdown language as implemented in the Redcarpet project, hosted on GitHub [2]. Redcarpet is a popular implementation of the language,

used by many other projects, in particular by the GitHub itself. Figure 4 shows the excerpt from the function `parse_table_header`, which can be found in the file `markdown.c`.

The function considered in the example parses the header of a table in the Markdown format. The first line of the header specifies column titles; they are separated by pipe symbols ('`|`'); the first pipe is optional. Thus, the number of pipes defines the number of columns in the table. The second line describes the alignment for each column, and should contain the same number of columns; in between each pair of pipes there should be at least three dash ('`-`') or colon ('`:`') symbols. A colon on the left or on the right side of the dashes defines left or right alignment; colons on both sides mean centered text. Thus, the two lines "`|One|Two|Three|`" and "`|:--|:--:|--:|`" specify three columns which are left-, center-, and right-aligned. Replacing the second line with either "`|:-|:--:|--:|`" or "`|:--|:--:|`" would result in the ill-formed input: the former doesn't contain enough dashes in the first column, while the latter doesn't specify the format for the last column.

Suppose, we are interested in the symbolic testing of the parser implementation; in particular, we want to cover all branches in the code for a reasonably long input. For that we postulate that the first input line contains at least n columns (we add the condition `assert(col>=n)` after line 19).

Now, consider the last conditional statement at line 19. The `if` branch is satisfied by an empty second input line; and indeed, such concolic testers as CREST can easily cover it. The `else` branch, however, poses serious problems. In order to cover it, a well-formed input that respects all constraints should be generated; in particular the smallest length of such input, e.g., for n equal to 3, is 17. The huge number of combinations to test exceeds the capabilities of the otherwise very efficient concolic tester: for $n = 2$ CREST needs 800 seconds to generate a test, and for $n = 3$ it is not able to finish within 3 hours.

Let us now examine the encoding of the implementation semantics in Array Folds Logic. The AFL assertions are shown in Figure 4 intertwined with the source code: they encode the semantics of the preceding code lines in the SSA form. To shorten the presentation we use the following conventions: variables `i`, `a`, `pipes`, `end`, `col`, and `dashes` are represented by (SSA-indexed) logical variables i, a, p, e, c, and d respectively; characters '`\n`', '`|`', '`:`', '`-`', and '`+`' by logical constants N, P, C, D, and A respectively; finally, the subscript denotes the SSA index of a variable.

The Presburger constraints such as those after line 2 are standard and we do not elaborate on them here. The first AFL-specific annotation goes after line 4: it directly reflects the loop semantics. The *fold* term encodes the computation of the number of pipes: they are computed in the counter $\mathbf{c}_1$, which gets its initial value equal to p_0, and its final value is equal to p_1. Similarly, array index $\mathbf{i}$ is initialized with i_0; and its final value is asserted to be equal to i_1. Both for counter $\mathbf{c}_1$ and for index $\mathbf{i}$ (which is a special type of a counter) their initial and final values can be both constant and symbolic: in fact, arbitrary Presburger terms are allowed.

```
1: static size_t parse_table_header(uint8_t *a, size_t size, ...)
2:   size_t i=0, pipes=0;
```
$$\left\{\ i_0 = 0 \ \wedge\ p_0 = 0\ \right\}$$
```
3:   while (i < size && a[i] != '\n')
4:     if (a[i++] == '|') pipes++;
```
$$\left\{\ \binom{i_1}{p_1} = \mathit{fold}_a\binom{i_0}{p_0}\left(\begin{smallmatrix}\mathbf{e}=P \Rightarrow \mathbf{c}_1\texttt{++}\\ \mathbf{e}\neq P \wedge \mathbf{e}\neq N \Rightarrow \mathit{skip}\end{smallmatrix}\right)\ \right\}$$
```
5:   if (a[0] == '|') pipes--;
```
$$\left\{\ \binom{*}{p_2} = \mathit{fold}_a\binom{0}{p_1}\left(\mathbf{i} = 0 \ \wedge\ \mathbf{e} = P \Rightarrow \mathbf{c}_1\texttt{--}\right)\ \right\}$$
```
6:  i++;
7:  if (i < size && a[i] == '|') i++;
```
$$\left\{\ i_2 = i_1 + 1 \ \wedge\ i_3 = \mathit{fold}_a(i_2)\left(\mathbf{i} = i_2 \ \wedge\ \mathbf{e} = P \ \Rightarrow\ \mathit{skip}\right)\ \right\}$$
```
8:   end = i;
9:   while (end < size && a[end] != '\n') end++;
```
$$\left\{\ e_0 = i_3 \ \wedge\ e1 = \mathit{fold}_a(e_0)\left(\mathbf{e} \neq N \ \Rightarrow\ \mathit{skip}\right)\ \right\}$$
```
10:  for (col = 0; col<pipes && i<end; ++col) {
11:    size_t dashes = 0;
```
$$\left\{\ c_0 = 0 \ \wedge\ c_0 < p_2 \ \wedge\ i_3 < e_1 \ \wedge\ d_0 = 0\ \right\}$$
```
12:    if (a[i] == ':') { i++; dashes++; column_data[col] |= ALIGN_L; }
```
$$\left\{\ \binom{i_4}{d_1} = \mathit{fold}_a\binom{i_3}{d_0}\left(\mathbf{i} = i_3 \ \wedge\ \mathbf{e} = C \Rightarrow \mathbf{c}_1\texttt{++}\right)\ \right\}$$
```
13:    while (i < end && a[i] == '-') { i++; dashes++; }
```
$$\left\{\ \binom{i_5}{d_2} = \mathit{fold}_a\binom{i_4}{d_1}\left(\mathbf{i} < e_1 \ \wedge\ \mathbf{e} = D \Rightarrow \mathbf{c}_1\texttt{++}\right)\ \right\}$$
```
14:    if (a[i] == ':') { i++; dashes++; column_data[col] |= ALIGN_R; }
```
$$\left\{\ \binom{i_6}{d_3} = \mathit{fold}_a\binom{i_5}{d_2}\left(\mathbf{i} = i_5 \ \wedge\ \mathbf{e} = C \Rightarrow \mathbf{c}_1\texttt{++}\right)\ \right\}$$
```
15:    if (i < end && a[i] != '|' && a[i] != '+') break;
16:    if (dashes < 3) break;
17:    i++;
```
$$\left\{\ (i_6 \geq \mathit{end}_1 \vee a[i_6] = P \vee a[i_6] = A) \ \wedge\ d_3 \geq 3 \ \wedge\ i_7 = i_6 + 1 \ \wedge\ c_1 = c_0 + 1\ \right\}$$
```
18:  }
19:  if (col < pipes) return 0;
```
$$\left\{\ c_1 \geq p_2\ \right\}$$

Fig. 4: An excerpt from the Redcarpet Markdown parser with AFL annotations

Notice that the loop at lines 3-4 is outside of the class of loops that can be accelerated by previous approaches. In particular, the difficulty here is the combination of the iteration over arrays with the branching structure inside the loop. On the contrary, AFL can summarize the loop in a concise logical formula.

The next conditional statement at line 5, takes care of the optional pipe at the beginning of the input. The annotation shown demonstrates that conditional statements are also easily represented by *fold* terms. In particular, here the function is folded over a starting from 0; the final index is unconstrained. The branch checks that the index is 0 (to prevent going further over the array), and that the symbol at this position is '|'. Counter $\mathbf{c}_1$ is decremented only if these two conditions are met; otherwise, the *fold* terminates. An equivalent encoding using only array reads is possible: $(a[0] = P \wedge p_2 = p_1 - 1) \vee (a[0] \neq P \wedge p_2 = p_1)$, but this encoding involves a disjunction.

The other program statements of the motivating example are encoded in a similar fashion. The encoding shown is for one unfolding of the `for` loop at line 10; several unfoldings are encoded similarly. We have checked the resulting proof obligations with our solver for AFL formulas, called AFolder; it can discharge them and generate the required test input in less than 2 minutes for $n = 3$.

4 Complexity

A counter machine is a finite automaton extended by a vector $\boldsymbol{\eta} = (\eta_1, \ldots, \eta_k)$ of k counters. Every counter in $\boldsymbol{\eta}$ stores a non-negative integer, and a counter machine can compare it to constant, and increment/decrease its value by a constant. For the formal definition of counter machines consult, e.g., [25].

We extend counter machines to symbolic counter machines (SCMs), which accept sequences (arrays) of integers. We denote the symbolic value of an array cell by a special integer variable $x_{\mathbf{e}}$. Let X be a set of integer variables, where $x_{\mathbf{e}} \notin X$. An *atomic input constraint* is of the form $x_{\mathbf{e}} \approx c$ or $x_{\mathbf{e}} \approx x$, where $c \in \mathbb{N}$, $x \in X$, and $\approx \in \{<, \leq, >, \geq, =, \neq\}$. Similarly, an *atomic counter constraint* is a formula of the form $\eta_i \approx c$ or $\eta_i \approx x$. An *input constraint* (resp. a *counter constraint*) is either a conjunction of $n \geq 1$ atomic input constraints (resp. atomic counter constraints), or the formula *true*. We denote by $\mathsf{IC}(X)$ (resp. $\mathsf{CC}_k(X)$) the set of all input constraints (resp. counter constraints with counters not greater than k) over variables in X.

Definition 1. *A* symbolic k-counter machine *is a tuple* $\mathcal{M} = (\boldsymbol{\eta}, X, Q, \delta, q^{\text{init}})$, *where:*

- $\boldsymbol{\eta} = (\eta_1, \ldots, \eta_k)$ *is a vector of* k *counter variables,*
- X *is a finite set of integer variables,*
- Q *is a finite set of states,*
- $\delta \subseteq Q \times \mathsf{CC}_k(X) \times \mathsf{IC}(X) \times Q \times \mathbb{Z}^k$ *is a transition relation,*
- $q^{\text{init}} \in Q$ *is the initial state.*

A transition $(q_1, \alpha, \beta, q_2, \boldsymbol{\kappa}) \in \delta$ moves the SCM from state q_1 to q_2 if the counters satisfy the constraint α and the inspected array cell satisfies β; the counters are incremented by $\boldsymbol{\kappa}$, and the machine moves to the next cell. A machine is called *deterministic* if δ is functional. A counter machine makes a *reversal* if it makes an alternation between non-increasing and non-decreasing some counter. A machine is *reversal-bounded* if there exists a constant $c \geq 0$ such that on all accepting runs every counter makes at most c reversal.

Definition 2. *We define the translation of a functional constant f of sort FSort^m, occurring in a formula ϕ, as an SCM $\mathcal{M}(f) = (\boldsymbol{\eta}, X, Q, \delta, q^{\text{init}})$. Let $G = \langle S, E, \gamma \rangle$ be the edge-labeled graph for f as defined in Section 2.1. Then $\boldsymbol{\eta} = \{\mathbf{i}, \mathbf{c}_1, \ldots, \mathbf{c}_{m-1}\}$, X are fresh free variables for each integer term T in f, $Q = S$, $q^{\text{init}} = 0$, and for each edge $(s_1, s_2) \in E$, δ contains a transition from s_1 to s_2 labeled with a conjunction of all constraints labeling the edge. For each integer term T in f and the corresponding variable $x \in X$, we replace T by x in f, and add the assertion $(x = T)$ as a conjunction at the outermost level of ϕ. Due to the constraint on G, we have that $\mathcal{M}(f)$ is reversal-bounded.*

Thus, we can translate a *fold* term into an SCM. A *parallel composition* of SCMs captures the scenario when several *fold*s operate over the same array.

Definition 3. *The parallel composition (product) of two SCMs $\mathcal{M}_1$ and $\mathcal{M}_2$, where $\mathcal{M}_i = (\boldsymbol{\eta}_i, X_i, Q_i, \delta_i, q_i^{\text{init}})$, is an SCM $\mathcal{M} = (\boldsymbol{\eta}, X, Q, \delta, q^{\text{init}})$ such that:*

- $\boldsymbol{\eta} = \boldsymbol{\eta}_1 \boldsymbol{\eta}_2$,
- $X = X_1 \cup X_2$,
- $Q = Q_1 \times Q_2$,
- *for each pair of transitions $(q_i, \alpha_i, \beta_i, p_i, \boldsymbol{w}_i) \in \delta_i$, where $i = 1..2$, there is the transition $\big((q_1, q_2), \alpha_1 \wedge \alpha_2, \beta_1 \wedge \beta_2, (p_1, p_2), \boldsymbol{w}_1 \boldsymbol{w}_2\big) \in \delta$, which are the only transitions in δ,*
- $q^{\text{init}} = (q_1^{\text{init}}, q_2^{\text{init}})$.

One of the fundamental questions that can be asked about a logic concerns the size of its models. The following lemma shows that models of bounded size are enough to check satisfiability of an AFL formula.

Lemma 1 (Small model property). *There exists a constant $c \in \mathbb{N}$, such that an AFL formula ϕ is satisfiable iff there exists a model σ such that a) for each integer variable x in ϕ, σ maps x to an integer $\leq 2^{|\phi|^c}$, and b) for each array variable in ϕ, σ maps the variable to a sequence of $\leq 2^{|\phi|^c}$ integers, where each integer is $\leq 2^{|\phi|^c}$.*

Proof (sketch; see the appendix for the full proof). One direction of the proof is trivial. For the other direction, assume that ϕ has a model σ. We construct a formula ψ that is a conjunction of all atomic formulas of ϕ: in positive polarity if σ satisfies the atomic formula, and in negative otherwise. Let $s = |\psi|$, and note that $s \leq 3|\phi|$. We observe that 1) σ is a model of ψ, 2) every model of ψ

is a model of ϕ. In the remaining part of the proof we show that ψ has a small model, and as a consequence so does ϕ.

Let a be some array in ψ. We translate each *fold* term over a to an SCM $\mathcal{M}_j$ as in Definition 2; let SCM $\mathcal{M}$ be the product of all $\mathcal{M}_j$. We extend the technique of [19] to show that there exist a sufficiently short run of $\mathcal{M}$. Under the interpretation σ, all variables in counter constraints become constants. Let $\boldsymbol{c} = (c_1, \ldots c_n)$ be a non-decreasing vector of constants that appear in the counter constraints of $\mathcal{M}$ after fixing σ. Vector $\boldsymbol{c}$ gives rise to the set of regions

$$\mathcal{R} = \{[0, c_1], [c_1, c_1], [c_1 + 1, c_2 - 1], [c_2, c_2], \ldots, [c_l, \infty]\}.$$

The size of $\mathcal{R}$ is at most $2\dim(\boldsymbol{c}) + 1 \leq 3s$. A mode of $\mathcal{M}$ is a tuple in $\mathcal{R}^k$ that describes the region of each counter. Let us observe that each counter can traverse at most $|\mathcal{R}|$ modes before it makes an additional reversal. Thus, $\mathcal{M}$ in any run can traverse at most $max = r \cdot k \cdot |\mathcal{R}| \leq \mathcal{O}(s^3)$ different modes.

We take some accepting run Tr of $\mathcal{M}$ that traverses at most max modes, and partition sequences of transition in Tr into equivalence classes. We create an integer linear program LP that encodes an accepting run of $\mathcal{M}$ that traverses at most max modes, as well as all non-fold constraints of ψ. The variables of LP correspond to 1) the integer variables of ψ, 2) the counter values of $\mathcal{M}$, 3) the number of times sequences from each equivalence class are taken, and 4) the solutions to each input constraint of $\mathcal{M}$.

We show that LP has a solution p, where each variable is at most $\leq 2^{|\phi|^c}$, for a fixed c. We use p to construct a small model for ψ. From p we immediately get interpretation for integer variables of ψ. Solution p implies that there is an accepting run of $\mathcal{M}$ of length at most $\leq 2^{|\phi|^c}$, which also gives a bound on the length of the input array. Finally, for every array cell we may use a solution to the specific input constraint. □

As a consequence of Lemma 1 we obtain a result on the complexity of AFL satisfiability checking.

Theorem 1. *The satisfiability problem of AFL is* **PSPACE***-complete.*

Proof. Membership. By Lemma 1, if an AFL formula ϕ is satisfiable, then it has a model where integer variables have value $\leq 2^{|\phi|^c}$, and arrays have length $\leq 2^{|\phi|^c}$, and where each array cell stores a number $\leq 2^{|\phi|^c}$. A non-deterministic Turing machine can use a polynomial number of bits to: 1) guess the value of integer variables and store them using $|\phi|^c$ bits, 2) guess one-by-one the value of at most $2^{|\phi|^c}$ array cells, and simulate the *fold*s. The Turing machine needs $|\phi|^c$ bits for counting the number of simulated cells. The maximum constant used in a counter increment can be at most $2^{|\phi|}$. Then, the maximal value a *fold* counter can store after traversing the array is at most $2^{|\phi|^{c+1}}$, therefore polynomial space is also sufficient to simulate *fold* counters.

Hardness. We reduce from the emptiness problem for intersection of deterministic finite automata, which is **PSPACE**-complete [27]. We are given a sequence $A_1, \ldots, A_n$ of deterministic finite automata, where each automaton A_i

accepts the language $\mathcal{L}(A_i)$. The problem is to decide whether $\bigcap_{i=1}^{n} \mathcal{L}(A_i) \neq \emptyset$. We simulate automata A_i with a *fold* expression $fold_a^i$ over a single counter, where input constraints correspond to the alphabet symbols of the automata. The expression $fold_a^i$ returns an even number on array a if and only if the interpretation of a represents a word in $\mathcal{L}(A_i)$. To check emptiness of the automata intersection, it is enough to check whether there exists an array such that all folds $fold_a^1, \ldots, fold_a^n$ return an even number. The reduction can be done in polynomial time. □

4.1 Undecidable Extensions

We show that two natural extensions to our logic lead to undecidability.

Theorem 2. *Array Folds Logic with an $\exists^*\forall^*$ quantifier prefix is undecidable.*

Proof. We prove by a reduction from Hilbert's Tenth Problem [30]; since addition is already in the logic, we only show how to encode multiplication. The following $\exists^*\forall^*$ AFL formula has a model iff array a is a repetition of z segments, and each segment is of length y and has the shape 00...01; thus, it asserts that $x = y \cdot z$:

$$|a| = x \ \wedge \ fold_a \binom{0}{0} \binom{\mathbf{e}{=}0 \Rightarrow skip}{\mathbf{e}{=}1 \Rightarrow \mathbf{c}_1\mathbf{{+}{+}}} = \binom{|a|}{z} \ \wedge$$
$$\forall j \,.\, 0 \leq j < |a| \implies fold_a \binom{j}{0} \binom{\mathbf{i} \leq j{+}y \wedge \mathbf{e}{=}0 \Rightarrow skip}{\mathbf{i} \leq j{+}y \wedge \mathbf{e}{=}1 \Rightarrow \mathbf{c}_1\mathbf{{+}{+}}} = \binom{*}{1} \qquad \square$$

In [11], the following is proved about the theory of concatenation:

Theorem 3 ([11], Corollary 4; see also [16], Proposition 1). *Solvability of equations in the theory $\langle \{1,2\}^*, e, \circ, Lg_1, Lg_2 \rangle$, where $Lg_p(x) \equiv \{y \in p^* \mid y \text{ has the same number of } p\text{'s as } x\}$, is undecidable.*

Corollary 1. *Extension of AFL with concatenation operator $\circ$ is undecidable.*

Proof. For an array x, we can define another array $Lg_1(x)$ in AFL as follows:

$$\binom{|x|}{|Lg_1(x)|} = fold_x \binom{0}{0} \binom{\mathbf{e}{=}1 \Rightarrow \mathbf{c}_1\mathbf{{+}{+}}}{\mathbf{e}{\neq}1 \Rightarrow skip} \ \wedge \ \binom{|Lg_1(x)|}{} = fold_{Lg_1(x)} \binom{0}{} \binom{\mathbf{e}{=}1 \Rightarrow skip}{} \qquad \square$$

5 Decision Procedure

In Section 4 we described how a non-deterministic Turing machine can decide AFL satisfiability in **PSPACE**. Now we present a deterministic procedure that translates AFL formulas to equisatisfiable quantifier-free Presburger formulas. As a consequence of the procedure, we show that under certain restrictions satisfiability of AFL is **NP**-complete.

Deterministic procedure We are given an AFL formula ϕ such that there are at most m *fold*s over each array; clearly m can be at most $|\phi|$. We translate ϕ to the quantifier-free Presburger formula $\psi = \psi_n \wedge \psi_e \wedge \psi_l$. For the procedure

we assume that there exists a fixed order $x_1 \leq \cdots \leq x_n$ on variables that appear in the counter constraints.

Formula ψ_n. The formula ψ_n is the part of ϕ that does not contain *fold*s.

Formula ψ_e. For an array a_j in ϕ, let $F_j = \{fold_a^1, \ldots, fold_a^m\}$ be the set of *fold*s in ϕ over a_i. We translate each $fold_a^i \in F_j$ to a symbolic counter machine $\mathcal{M}_j^i$. Each $\mathcal{M}_j^i$ has at most $|\phi|$ transitions, and the sum of the counters and the number of reversals among all $\mathcal{M}_j^i$ is at most $|\phi|$. Next, we construct the symbolic counter machine $\mathcal{M}_j$ as the product of all machines $\mathcal{M}_j^i$. The machine $\mathcal{M}_j$ has at most $k = |\phi|$ counters, $t = |\phi|^m$ transitions and makes at most $r = |\phi|$ reversals.

We translate the reachability problem of $\mathcal{M}_j$ to the quantifier-free Presburger formula ψ_e^j by applying an extension of the method described in [22]. In formula ψ_e^j, two configurations of $\mathcal{M}_j$ are described symbolically: initial ζ, and final ζ'. The formula ψ_e^j is satisfiable iff there is an array a_j such that $\mathcal{M}_j$ reaches ζ' from ζ on reading a_j. The formula ψ_e is the conjunction ψ_e^j for all arrays a_j.

The formula ψ_e^j consists of two parts $\psi_e^j = \psi_p^j \wedge \psi_c^j$. For simplicity we assume that the counter constrains of $\mathcal{M}_j$ are defined only over variables $\{x_1, \cdots, x_n\}$. By assumption, there is a fixed order $x_1 \leq \ldots \leq x_n$, which gives rise a to the set of $\leq 2|\phi| + 1$ regions $\mathcal{R} = \{[0, x_1], [x_1, x_1], [x_1 + 1, x_2 - 1], \cdots, [c_l, \infty]\}$. As an optimization, we construct regions separately for each counter, which allows us to obtain a tighter bound on the number of regions that need to be encoded.

Each counter may traverse at most $|\mathcal{R}|$ regions before it makes a reversal, so an accepting computation of $\mathcal{M}_j$ traverses at most $max = r \cdot k \cdot |\mathcal{R}| = \mathcal{O}(|\phi|^3)$ modes. We construct an NFA $\mathcal{A}_j$ by making max copies of the control-flow structure of $\mathcal{M}_j$. Every run of $\mathcal{A}_j$ gives a correct sequence of states in $\mathcal{M}_j$, but may violate counter constraints. By using the procedure of [34] we can encode the Parikh image of $\mathcal{A}_j$ as the formula ψ_j^p that is polynomial in the size of $\mathcal{A}$. Similar to [22], the formula ψ_c^j puts additional constraints on the Parikh image to ensure that by executing the transitions of $\mathcal{A}_j$ we obtain counter values that satisfy the counter constraints of $\mathcal{M}_j$.

The size ψ_e^j is of the order $\mathcal{O}(|\phi|^3 t) = \mathcal{O}(|\phi|^{m+3})$. The formula ψ_e is the conjunction of formulas ψ_e^j for each array a_j. There can be at most $|\phi|$ arrays, so the size of ψ_e is $\mathcal{O}(|\phi|^{m+4})$.

Formula ψ_l. Finally, formula ψ_l links the initial and final configurations in ψ_e to the variables in ψ_p.

Formula size. The size of the formula ψ is $\mathcal{O}(|\phi|^{m+4})$. By keeping m constant, the encoding size is polynomial in the size of the AFL formula ϕ.

Restricted fragment of AFL We write m-AFL for formulas that have at most m *fold* expressions per array. As a consequence of the deterministic decision procedure, restriction on m reduces the complexity of deciding satisfiability.

Lemma 2. *The m-AFL satisfiability problem, for a fixed m, is* **NP***-complete.*

Proof. Membership follows from the decision procedure above. For hardness observe that any quantifier-free Presburger formula is an 0-AFL formula. □

Model generation Given a Presburger encoding ψ of an AFL formula ϕ, we may use the solution to ψ to generate a model of ϕ. The solution to ψ immediately gives us interpretation for the integer variables in ϕ. To obtain an interpretation for the array variables in ϕ, we observe that folds are implicitly encoded in ψ as counter machines, and that the solution to ψ describes the Parikh vector for each machine. We use the method of [34] to get a concrete sequence of transitions in each counter machine that produces the specific Parikh vector. We construct a multigraph by repeating each transition in $\mathcal{A}_j$ according to its Parikh image, and then find an Eulerian path in the multigraph. From the sequence of transitions in counter machines, and the interpretation of input constraints in ψ we obtain an interpretation for the arrays in ϕ.

6 Experiments

We implemented the decision procedure described in Section 5 in a prototype tool AFolder; the tool is available at [13]. The tool is written in `C++` and uses Z3 [14] as the solver for Presburger formulas. We evaluated our decision procedure on a number of testing and verification tasks described below.

The experimental results are shown in Table 3; all experiments were performed on a Ubuntu-14.04 64-bit machine running on an Intel Core i5-2540M CPU of 2.60GHz. For every example we report the size $|\phi|$ of the AFL formula measured as "the number of logical operators" + "the number of branches in folds." The table also shows the number of *fold* expressions in a formula, and the maximum number of folds per array (MFPA). Next, we report the time for translating the problem to a Presburger formula, the time for solving the formula, and whether the formula is satisfiable. If this is the case, we report the length of a satisfying array generated by our tool; in case of several arrays, we show the longest.

Markdown This program is described in Section 3. The experiments are parametrized by the required number n of columns in the input.

perf_bench_numa This example is part of a benchmark program for non-uniform memory access (NUMA) [1]. The program maintains a list of threads, and for each thread a separate array of size 100 that describes processors assigned to the thread. The data is processed in a nested loop: the outer loop iterates over threads, and the inner loop counts the number of assigned processors. The outer loop also maintains the minimum, and maximum number of processors assigned to any thread. We model a testing scenario like in Section 3, where a symbolic execution tool unrolls the outer loop n times, and the inner loop is summarized by a fold expression. The testing goal is to provide a valid processor mapping such that each thread is assigned to exactly one processor. In Table 3 we show results for this benchmark parametrized by the number n of threads. The example scales well, since there a single fold per each processor array (see Lemma 2).

SV-COMP Examples "standard_minInArray" to "standard_vararg" are taken from the SV-COMP benchmarks suite [6]. They model simple verification prob-

Table 3: Experimental results for AFolder.

Example	$\|\phi\|$	folds	MFPA	transl. time	solving time	result	array length
Markdown(1)	62	6	3	$< 1s$	$< 1s$	sat	8
Markdown(2)	69	7	4	$1s$	$< 1s$	sat	14
Markdown(3)	76	8	5	$1.3s$	79s	sat	17
perf_bench_numa(10)	93	10	1	< 1s	< 1s	sat	100
perf_bench_numa(20)	183	20	1	< 1s	< 1s	sat	100
perf_bench_numa(40)	363	40	1	< 1s	< 1s	sat	100
standard_minInArray	10	3	3	< 1s	< 1s	unsat	-
linear_sea.ch_true	13	3	3	< 1s	< 1s	unsat	-
array_call3	11	2	3	< 1s	< 1s	unsat	-
standard_sentinel	14	3	3	< 1s	< 1s	unsat	-
standard_find	11	3	3	< 1s	< 1s	unsat	-
standard_vararg	11	3	3	< 1s	< 1s	unsat	-
histogram(8)	58	8	8	< 1s	1.3s	sat	9
histogram(9)	65	9	9	< 1s	6.9s	sat	10
histogram(10)	72	10	10	2s	55s	sat	11
histogram(11)	79	11	11	8s	368s	sat	12
histogram_unsat(11)	80	11	11	9s	19s	unsat	-

lems for loops, such as finding the position of an element in array, finding the minimum, or counting the number of positive elements. We model these programs as formulas that are unsatisfiable if the program is safe. Although the programs are simple, most verification tools competing in SV-COMP fail to prove their safety.

histogram We performed experiments on the histogram example in Section 2.3, parametrized by the number of range values. We observe that solving time grows rapidly with the number of folds. Example "histogram_unsat" is an unsatisfiable variation that requires two different counts in the same range.

7 Conclusion and Future Work

We presented Array Folds Logic (AFL), which extends the quantifier-free theory of arrays with *folding*, a well-known concept from functional languages. The extension allows us to express counting properties, occurring frequently in real-life programs. Additionally, AFL is able to concisely summarize loops with internal branching and counting over arrays. We have analyzed the complexity of satisfiability checking for AFL formulas, and presented an efficient decision procedure via an encoding to the quantifier-free Presburger arithmetic. Finally, we have implemented a tool called AFolder, which efficiently discharges AFL proof obligations, and demonstrated its practical applicability on numerous examples.

For the future work, we plan to investigate possible combinations with other decidable fragments of the theory of arrays (to allow some restricted form of quantifier alternation). We also plan to automate the generation of proof obligations and the summarization of loops, and want to improve the efficiency of our decision procedure by implementing suitable optimizations and heuristics.

References

1. Benchmark for non-uniform memory access (NUMA). `http://lxr.free-electrons.com/source/tools/perf/bench/numa.c?v=3.9`. Accessed: 2016-01-29.
2. Markdown parser of the Redcarpet project. `https://github.com/vmg/redcarpet/blob/master/ext/redcarpet/markdown.c`. Accessed: 2016-01-20.
3. F. Alberti, S. Ghilardi, and N. Sharygina. Booster: An acceleration-based verification framework for array programs. In F. Cassez and J. Raskin, editors, *Automated Technology for Verification and Analysis - 12th International Symposium, ATVA 2014, Sydney, NSW, Australia, November 3-7, 2014, Proceedings*, volume 8837 of *Lecture Notes in Computer Science*, pages 18–23. Springer, 2014.
4. F. Alberti, S. Ghilardi, and N. Sharygina. Decision procedures for flat array properties. *J. Autom. Reasoning*, 54(4):327–352, 2015.
5. F. Alberti, S. Ghilardi, and N. Sharygina. A new acceleration-based combination framework for array properties. In C. Lutz and S. Ranise, editors, *Frontiers of Combining Systems - 10th International Symposium, FroCoS 2015, Wroclaw, Poland, September 21-24, 2015. Proceedings*, volume 9322 of *Lecture Notes in Computer Science*, pages 169–185. Springer, 2015.
6. D. Beyer. Software verification and verifiable witnesses - (report on SV-COMP 2015). In *TACAS*, pages 401–416, 2015.
7. M. Bozga, P. Habermehl, R. Iosif, F. Konecný, and T. Vojnar. Automatic verification of integer array programs. In A. Bouajjani and O. Maler, editors, *Computer Aided Verification, 21st International Conference, CAV 2009, Grenoble, France, June 26 - July 2, 2009. Proceedings*, volume 5643 of *Lecture Notes in Computer Science*, pages 157–172. Springer, 2009.
8. M. Bozga, R. Iosif, and F. Konecný. Fast acceleration of ultimately periodic relations. In T. Touili, B. Cook, and P. Jackson, editors, *Computer Aided Verification, 22nd International Conference, CAV 2010, Edinburgh, UK, July 15-19, 2010. Proceedings*, volume 6174 of *Lecture Notes in Computer Science*, pages 227–242. Springer, 2010.
9. M. Bozga, R. Iosif, F. Konecný, and T. Vojnar. Tool demonstration of the FLATA counter automata toolset. In A. Voronkov, L. Kovács, and N. Bjørner, editors, *Second International Workshop on Invariant Generation, WING 2009, York, UK, March 29, 2009 and Third International Workshop on Invariant Generation, WING 2010, Edinburgh, UK, July 21, 2010*, volume 1 of *EPiC Series*, page 75. EasyChair, 2010.
10. A. R. Bradley, Z. Manna, and H. B. Sipma. What's decidable about arrays? In E. A. Emerson and K. S. Namjoshi, editors, *VMCAI*, volume 3855 of *LNCS*, pages 427–442. Springer, 2006.
11. J. R. Büchi and S. Senger. Definability in the existential theory of concatenation and undecidable extensions of this theory. *Math. Log. Q.*, 34(4):337–342, 1988.
12. H. Comon and Y. Jurski. Multiple counters automata, safety analysis and presburger arithmetic. In A. J. Hu and M. Y. Vardi, editors, *Computer Aided Verification, 10th International Conference, CAV '98, Vancouver, BC, Canada, June 28 - July 2, 1998, Proceedings*, volume 1427 of *Lecture Notes in Computer Science*, pages 268–279. Springer, 1998.
13. P. Daca. AFolder. `https://github.com/pdaca/AFolder`.
14. L. M. de Moura and N. Bjørner. Z3: An efficient SMT solver. In *TACAS*, 2008.

15. L. M. de Moura and N. Bjørner. Generalized, efficient array decision procedures. In *Proceedings of 9th International Conference on Formal Methods in Computer-Aided Design, FMCAD 2009, 15-18 November 2009, Austin, Texas, USA*, pages 45–52. IEEE, 2009.
16. C. A. Furia. What's decidable about sequences? In A. Bouajjani and W. Chin, editors, *Automated Technology for Verification and Analysis - 8th International Symposium, ATVA 2010, Singapore, September 21-24, 2010. Proceedings*, volume 6252 of *Lecture Notes in Computer Science*, pages 128–142. Springer, 2010.
17. S. Ghilardi, E. Nicolini, S. Ranise, and D. Zucchelli. Decision procedures for extensions of the theory of arrays. *Ann. Math. Artif. Intell.*, 50(3-4):231–254, 2007.
18. P. Godefroid and D. Luchaup. Automatic partial loop summarization in dynamic test generation. In *Proceedings of the 20th International Symposium on Software Testing and Analysis, ISSTA 2011, Toronto, ON, Canada, July 17-21, 2011*, pages 23–33, 2011.
19. E. M. Gurari and O. H. Ibarra. The complexity of decision problems for finite-turn multicounter machines. In *Automata, Languages and Programming*, pages 495–505, 1981.
20. P. Habermehl, R. Iosif, and T. Vojnar. A logic of singly indexed arrays. In I. Cervesato, H. Veith, and A. Voronkov, editors, *LPAR*, volume 5330 of *LNCS*, pages 558–573. Springer, 2008.
21. P. Habermehl, R. Iosif, and T. Vojnar. What else is decidable about integer arrays? In R. M. Amadio, editor, *Foundations of Software Science and Computational Structures, 11th International Conference, FOSSACS 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29 - April 6, 2008. Proceedings*, volume 4962 of *Lecture Notes in Computer Science*, pages 474–489. Springer, 2008.
22. M. Hague and A. W. Lin. Model checking recursive programs with numeric data types. In *CAV*, pages 743–759, 2011.
23. H. Hojjat, R. Iosif, F. Konecný, V. Kuncak, and P. Rümmer. Accelerating interpolants. In S. Chakraborty and M. Mukund, editors, *Automated Technology for Verification and Analysis - 10th International Symposium, ATVA 2012, Thiruvananthapuram, India, October 3-6, 2012. Proceedings*, volume 7561 of *Lecture Notes in Computer Science*, pages 187–202. Springer, 2012.
24. O. H. Ibarra. Reversal-bounded multicounter machines and their decision problems. *J. ACM*, 25(1):116–133, 1978.
25. O. H. Ibarra, J. Su, Z. Dang, T. Bultan, and R. A. Kemmerer. Counter machines and verification problems. *Theor. Comput. Sci.*, 289(1):165–189, 2002.
26. J. Knoop, L. Kovács, and J. Zwirchmayr. Symbolic loop bound computation for WCET analysis. In *Perspectives of Systems Informatics - 8th International Andrei Ershov Memorial Conference, PSI 2011, Novosibirsk, Russia, June 27-July 1, 2011, Revised Selected Papers*, pages 227–242, 2011.
27. D. Kozen. Lower bounds for natural proof systems. In *FOCS*, pages 254–266, 1977.
28. A. W. Lin and P. Barceló. String solving with word equations and transducers: towards a logic for analysing mutation XSS. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, pages 123–136.
29. G. S. Makanin. The problem of solvability of equations in a free semigroup. *Sbornik: Mathematics*, 32(2):129–198, 1977.
30. J. V. Matijasevič. The Diophantineness of enumerable sets. *Dokl. Akad. Nauk SSSR*, 191:279–282, 1970.

31. J. McCarthy. Towards a mathematical science of computation. In *IFIP Congress*, pages 21–28, 1962.
32. C. H. Papadimitriou. On the complexity of integer programming. *J. ACM*, 28(4):765–768, 1981.
33. P. Saxena, P. Poosankam, S. McCamant, and D. Song. Loop-extended symbolic execution on binary programs. In *Proceedings of the Eighteenth International Symposium on Software Testing and Analysis, ISSTA 2009, Chicago, IL, USA, July 19-23, 2009*, pages 225–236, 2009.
34. H. Seidl, T. Schwentick, A. Muscholl, and P. Habermehl. Counting in trees for free. In *Automata, Languages and Programming*, pages 1136–1149. Springer, 2004.
35. N. Suzuki and D. Jefferson. Verification decidability of presburger array programs. *J. ACM*, 27(1):191–205, 1980.
36. M. Zhou, F. He, B.-Y. Wang, M. Gu, and J. Sun. Array theory of bounded elements and its applications. *Journal of automated reasoning*, 52(4):379–405, 2014.

A Proof of Lemma 1

Proof. One direction of the proof is trivial.

For the other direction, assume that ϕ has a model σ. Let X be the set of variables in ϕ. W.l.o.g. we assume that all folds are of the form

$$\binom{out_1}{\cdots}{out_n} \;=\; fold_a \binom{in_1}{\cdots}{in_n} F^m$$

where $out_1, \ldots, out_n, in_1, \ldots, in_n$ are integer variables.

From the model σ of ϕ we build a conjunction ψ of literals in the following way: for every atomic formula γ of the form $T = T, V^m = V^m$, or $T \leq T$ in ϕ we add a conjunct γ to ψ if σ satisfies γ, and we add a conjunct $\neg\gamma$ otherwise. Observe that σ is a model of ψ and every model of ψ is a model of ϕ. In the remaining part of the proof we show that ψ has a small model.

Let $s = |\psi|$, and note that $s \leq 3|\phi|$. Moreover, we write $\psi = \psi_f \wedge \psi_{nf}$, where ψ_f contains only literals with folds, and ψ_{nf} contains only literals without folds.

Let us assume that folds are over the same array a; we will later deal with this restriction. Let $F = \{fold_a^1, \ldots, fold_a^n\}$ be a set of folds in ψ over the array a. We translate each fold $fold_a^i \in F$ to a symbolic counter machine $\mathcal{M}_i$. Each $\mathcal{M}_i$ has at most s transitions, and the sum of counters and reversal among all $\mathcal{M}_i$ is at most s. Next, we create the product $\mathcal{M} = (\boldsymbol{\eta}, Q, \Sigma, \delta, q^{\text{init}})$ of all the folds in F. The product counter machine $\mathcal{M}$ has at most $k = s$ counters, s^s states and transitions, and makes at most $r = s$ reversals.

In the following part of the proof, we extend the technique of [19] to show that there exist a sufficiently short run of $\mathcal{M}$. Under the interpretation σ, all variables in counter constraints become constants. Let $\boldsymbol{c} = (c_1, \ldots c_n)$ be a non-decreasing vector of constants that appear in the counter constraints of $\mathcal{M}$ after fixing σ. Vector $\boldsymbol{c}$ gives rise to a set of regions

$$\mathcal{R} = \{[0, c_1], [c_1, c_1], [c_1 + 1, c_2 - 1], [c_2, c_2], \ldots, [c_l, \infty]\}.$$

The size of $\mathcal{R}$ is at most $2\dim(\boldsymbol{c}) + 1 \leq 3s$. A mode of $\mathcal{M}$ is a tuple in $\mathcal{R}^k$ that describes the region of each counter. Let us observe that each counter can traverse at most $|\mathcal{R}|$ modes before it makes an additional reversal. Thus, $\mathcal{M}$ in any run can traverse at most $max = r \cdot k \cdot |\mathcal{R}| \leq \mathcal{O}(s^3)$ different modes.

Let $\boldsymbol{w}^{in}, \boldsymbol{w}^{out} \in \mathbb{N}^k$ be the vectors of the initial and final values of counters of $\mathcal{M}$; these vectors are given by interpretation σ of the initial and output variables of folds. We know that $\mathcal{M}$ on input $\sigma(a)$ has a run $Tr = \delta_1, \ldots, \delta_n$, where $\delta_i \in \delta$, from an initial configuration $\zeta_0 = (q^{\text{init}}, \boldsymbol{w}_{in})$ to a final configuration $\zeta_n = (\cdot, \boldsymbol{w}_{out})$.

We partition Tr into sub-sequences $Tr_1, \ldots, Tr_{max}$, such that all transitions in Tr_i are fired from a configuration in mode i, and for $i < max$ the last transition in Tr_i leads to a configuration in mode $i + 1$. Let us look into some sub-run $Tr_i = \delta_l \ldots \delta_{l'}$. For each transition $\delta \in \delta$ we mark one occurrence of δ in Tr_i, provided that such transition occurs. In this way, we mark at most $|\delta| \leq s^s$ transitions in Tr_i.

Next, we identify sub-sequences $\rho = \delta_m, \ldots, \delta_n$ of Tr_i, such that 1) $m < n$, 2) $\delta_m = \delta_n$, 3) all transitions in ρ are unmarked, and 4) all transitions $\delta_{m+1}, \ldots, \delta_{n-1}$ are distinct. Observe that by deleting transitions $\delta_2, \ldots, \delta_l$ from Tr_i we obtain a valid sequence of transition that ends in the same state, but may lead to different counter values. Let $\bar{Tr}_i$ be the sequence of transitions that results by repeatedly deleting all such sequences from Tr_i, and let S_i be the multi-set of sequences deleted in such way. Since marked configurations remain in $\bar{Tr}_i$ (and there are at most $|\delta|$ of them), and there can be at most $|\delta|$ remaining transition between any two marked configuration in $\bar{Tr}_i$, therefore $\bar{Tr}_i$ has at most $|\delta|(1 + |\delta|) = \mathcal{O}(s^{3s})$ transitions.

For each S_i we define an equivalence relation $S_i^{=}$ on the deleted transitions as follows: two deleted sequences are equivalent if they 1) have the same starting transitions, 2) have the same end transitions, and 3) add/subtract the same value for each counter. Note that in Tr_i a deleted sub-sequence can be substituted by an equivalent one, without changing the final configuration, i.e. the state and the value of counters remain the same. Let u be the maximum constant which is added/subtracted in a counter update; u is given in binary, so $u \leq 2^s$. There may be at most $|\delta|$ transitions in a deleted sequence, so the total net effect on a single counter may be at most $u|\delta| = s^s 2^s$. The number of equivalence classes in $S_i^{=}$ is $|S_i^{+}| \leq u|\delta|^2 \leq s^{3s}2^s \leq \mathcal{O}(s^{4s})$.

We construct an integer linear program LP, which expresses solutions to ψ that lead to some run Tr^* of $\mathcal{M}$ such that 1) Tr^* goes through the same modes as Tr, 2) Tr^* has the same net effect on each counter, 3) Tr^* is possibly shorter than Tr. The linear program has five parts $LP = LP_1 \wedge \ldots \wedge LP_5$.

$\mathbf{LP_1}$ In LP_1 we specify constraints for the counter values. For every counter $1 \leq j \leq k$ and mode $1 \leq i \leq max$ we create variables $w_{i,j}$ and $z_{i,j}$ which describe the value of counter j at the start and the end of mode i. Thus, $z_{max,j}$ is the value of counter j after executing Tr^*, and $w_{1,j}$ is the initial value of the computation.

For $1 \leq j \leq k, 1 \leq i \leq max$, part LP_1 contains the following equations

$$z_{i,j} = w_{i,j} + \bar{b}_{i,j} + \sum_{m=1}^{|S_i^{=}|} b_{i,j,m}\ y_{i,m},$$

where $\bar{b}_{i,j} \in \mathbb{N}$ is the net effect of $\bar{Tr}_i$ on counter j, $b_{i,j,m} \in \mathbb{N}$ is the net effect of the sequence in equivalence class m on counter j, and variable $y_{i,m}$ denotes the number of times that the sequence from class m occurs in mode i.

Additionally, for $1 \leq i < max - 1$, we add to LP_1 the constraint:

$$w_{i+1,j} = z_{i,j} + b^{+}_{i,j},$$

where $b^{+}_{i,j}$ is the effect of the last transition in $\bar{Tr}_i$ on counter j. We denote by Z, W, Y the sets of variables z, w, y, respectively.

$\mathbf{LP_2}$ In LP_2 we specify the constrain for ψ_{nf}, which is the part of ψ that does not contain folds.We create a copy X' of all the variables in X. The linear

program contains a formula

$$LP_2 = \psi_{nf}[X/X'].$$

LP$_3$ We link the start and end value of counters with variables in X'. For a counter $1 \leq j \leq k$ we add the following constraints to LP_3

$$z_{max,j} = out'_k \wedge w_{i,j} = in'_k,$$

where out_k, in_k are the variables that specify in ψ the initial and output values of counter j of $\mathcal{M}$.

LP$_4$ W.l.o.g. we assume that in the folds of ψ counters are only compared to variables. As a result of our translation, the constants in counter constraints of $\mathcal{M}$ arise from interpretation σ of some variables in X. Thus, we add the constraint that counter values stay within their modes. Suppose in mode i, counter j is in the region specified by the interval $[x_k, x_l]$, where $x_k, x_l \in X$. We add the following constraints to LP_4:

$$x'_l \leq w_{i,j} \leq x'_k \quad \wedge \quad x'_l \leq z_{i,j} \leq x'_k.$$

LP$_5$ Finally, we need to ensure that all input constraints executed in Tr are satisfiable. Thus, let $\mathcal{I}$ be the set of input constraints executed in Tr, where $|\mathcal{I}| \leq s$. For every input constraint $\psi_i \in \mathcal{I}$ we create an new variable v_i and add to LP_5 a linear constraint over v^i that corresponds to ψ_i. We denote by V the set of v_i variables.

The linear program LP can be expressed in the form $\boldsymbol{A}\boldsymbol{x} \leq \boldsymbol{b}$, where $\boldsymbol{x}$ is a vector of variables in $X' \cup Z \cup W \cup Y \cup V$, and and $\boldsymbol{A}, \boldsymbol{b}$ are a matrix, and vector of integers. To determine the dimension of $\boldsymbol{x}$, observe that the cardinalities of X', Z, W, V are polynomial in s, while

$$|Y| \leq \sum_{i=1}^{max} |S_i^{=}| \leq \mathcal{O}(s^{k_0 s}),$$

where k_0 is a fixed constant. Thus, $\dim(\boldsymbol{x}) \leq \mathcal{O}(s^{k_0 s})$. Each of the constraints in $LP_1, \ldots, LP_2$ adds a number of constrains polynomial in s, so $\dim(\boldsymbol{b}) \leq \mathcal{O}(s^{k_1})$, where k_1 is a fixed constant. Observe that all entries in $\boldsymbol{A}$ and $\boldsymbol{b}$ are at most of the size 2^s – the largest number we can encode in ψ. As a result the size of $[\boldsymbol{A}\boldsymbol{b}]$ is at most $\mathcal{O}(s^{sk_2})$, where k_2 is a fixed constant.

We know that there exists a solution to $\boldsymbol{x}$ that satisfies the linear program: this solution can be obtained be assigning appropriate variables to the values of Tr and σ. By [32] we know that there also exists a solution p to $\boldsymbol{x}$ to assigns to every variable an element with absolute value is at most $2^{s^{k_3}}$, for some fixed constant k_3 (to satisfy the requirements of [32], we transforms LP program to the standard form $\boldsymbol{A'x'} = \boldsymbol{b'}, \boldsymbol{x} \geq 0$ – this can be done by doubling the number of constraints and variables. The solutions to $\boldsymbol{x'}$ can be mapped to the solutions in $\boldsymbol{x}$).

From the solution p we can construct a model σ' of ψ in the two steps. First, for integer variables in X, assign the values of X' under p. Clearly, each variable gets a value with a representation that is $\mathcal{O}(s^{k_3})$ bits.

Second, to get an assignment σ' to array a we create a computation Tr^* that has the same effect as Tr. To obtain Tr_i^*, in every reduced sub-sequence $\bar{Tr}_i$ place $p(y_{i,m})$ copies of a sequence in the equivalence class m of $S_i^=$ after the "marked transition rule identical to the last transition rule in the sequence" (see [19]). Then $Tr^* = Tr_1^*, \ldots, Tr_{max}^*$, and the length of Tr^* is in $\mathcal{O}(s^{sk_4})$, where k_4 is a fixed constant. The assignment to array a is a sequence of integers which satisfy the input constraints that run Tr^* reads. For each input constraint ψ_i read by Tr^*, we may use a solution p to v_i as a concrete value. The assignment to array a is of length $\mathcal{O}(s^{sk_4})$, and each element can be represented in at most $\mathcal{O}(s^{k_3})$ bits. To complete the proof, let us observe that there can be at most $|\psi|$ arrays in the formula ψ, so having multiple arrays does not change the order of growth of σ'. □